

Spin glass patch planting

Wenlong Wang,^{1,*} Salvatore Mandrà,^{2,3,†} and Helmut G. Katzgraber^{1,4,5}

¹*Department of Physics and Astronomy, Texas A&M University, College Station, Texas 77843-4242, USA*

²*NASA Ames Research Center Quantum Artificial Intelligence Laboratory (QuAIL), Mail Stop 269-1, 94035 Moffett Field CA*

³*Stinger Ghaffarian Technologies Inc., 7701 Greenbelt Rd., Suite 400, Greenbelt, MD 20770*

⁴*Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, New Mexico 87501, USA*

⁵*Applied Mathematics Research Centre, Coventry University, Coventry, CV1 5FB, England*

In this paper, we propose a patch planting method for creating *arbitrarily large* spin glass instances *with known ground states*. The scaling of the computational complexity of these instances with various block numbers and sizes is investigated and compared with random instances using population annealing Monte Carlo and the quantum annealing DW2X machine. The method can be useful for benchmarking tests for future generation quantum annealing machines, classical and quantum mechanical optimization algorithms.

PACS numbers: 75.50.Lk, 75.40.Mg, 05.50.+q, 64.60.-i

I. INTRODUCTION

Many optimization problems belong to the important class of NP-Hard problems, for which it is believed that no algorithms exist to solve them in polynomial time. Spin glass problems like the Edward-Anderson (EA) model [1] represent a sub-class of the NP-Hard class. Hence, not only they share the same computational hardness of all the other members of the class, but also any NP-Hard problem (or, in general, any problems belong the NP class) can be mapped onto a spin glass Hamiltonian with only a polynomial overhead. A number of heuristics, as well as exhaustive search methods, have been designed and developed to solve these problems as efficiently as possible. These methods including simulated annealing [2], parallel tempering [3–5], population annealing [6–9], genetic algorithms [10, 11], as well as DPLL and branch-and-cut [12] algorithms. Many of these optimization algorithms only use local move to eventually reach the ground state. However, in many cases, the use of cluster algorithms with nonlocal updates can greatly enhance the searching process when the energy landscape presents many metastable states with small overlap [13, 14]. In the last two decades, quantum heuristics have been also proposed as an alternative for classical heuristics, thanks to their potentiality to exploit quantum superposition and quantum tunneling. Among the quantum technologies, adiabatic quantum optimization (AQO) is one of the most used [15–29] and the more viable from the manufacturing point of view [30]. The state-of-art AQO hardware is manufactured by D-Wave System Inc., whose latest chip allows the quantum optimization of 1000 qubits. However, whether AQO can be more efficient than classical algorithms in certain problems is still controversial [31–33].

Given the importance to compare heuristics that can be so different from each other, it is mandatory to define a benchmark class with the properties to be i) representative of the hardness of a typical NP-Hard problem, ii) scalable for large systems and for which iii) the ground state is known *a priori*. Indeed, if it easy to find benchmark class that fulfill i) and ii), it becomes challenging to find problems with thousands of variables but for which the ground state is still known.

In order to fulfill all the properties which are expected for a good benchmark class, we introduce a new method called “patch & planting” (PP). The method consists in patching together smaller problems, for which the ground state can be numerically found, to form larger instance for which the ground state is still known but computationally harder to find. The paper is structures as follows: In Sec. II we introduce the PP model, showing how to construct instances with millions of qubits. In Sec. III, we numerically show that problems created with our method are indeed hard to solve. Results are obtained either numerically using simulated annealing or population annealing, or experimentally using the D-Wave 2X chip. Concluding remarks are finally stated in Sec. IV.

II. MODELS AND METHODS

A. Edwards-Anderson model

The Edward-Anderson (EA) model [1] is one of the most known and well studied spin glass problems, for which many properties are known either analytically or numerically. For example, it is known that EA has a finite spin glass temperature when three-dimensional lattice are considered. On the contrary, for two-dimensional lattices, it is expected that EA has a zero spin glass critical temperature. The EA is defined by the following spin Hamiltonian

$$H = - \sum_{\langle ij \rangle} J_{ij} S_i S_j - \sum_i h_i S_i, \quad (1)$$

*Electronic address: wenlong@physics.umass.edu

†Electronic address: salvatore.mandra@nasa.gov

where $S_i \in \{\pm 1\}$ are Ising spins and the first sum is over nearest neighbors on a d -dimensional lattice of linear size L . In this work, we only considered 3-dimensional lattices. For simplicity, all the local magnetic fields are set to zero, namely $h_i = 0$. Nevertheless, our method works with external fields as well. The couplings J_{ij} define the “disorder” and they are randomly chosen from a given distribution. Here, we considered independent and identically distributed random J_{ij} , extracted from a Gaussian distribution with zero mean and unitary variance. The set of the couplings also unequivocally define an “instance”.

Given the hardware limitation of the D-Wave quantum chips, instances for the D-Wave 2X have been created by planting and patching together unit cells following the two-dimensional architecture of the Chimera graph. Moreover, we used couplings randomly extracted from the Sidon set $\{\pm 5, \pm 6, \pm 7\}$. In both cases, we use free boundary conditions (FBC) for the blocks to plant larger instances. For the EA model, we also compare FBC with random instances with regular periodic boundary conditions (PBC).

B. Patch planting

As briefly described in the introduction, the “patch & planting” (PP) method works by following the simple scheme:

1. Solve some small problems, called blocks, using *free boundary conditions* and find their ground states.
2. For each block, choose a ground state.
3. Add couplings between blocks so that all the extra-block couplings are “satisfied” with respect to the chosen ground states.

Observe that the aforementioned scheme always succeeds since all the extra-couplings can be either positive or negative to accordingly match aligned or anti-aligned spins. Moreover, since all the extra-block couplings are satisfied, the ground state for the larger instance is necessarily the composition of all the ground states of the blocks. A schematic diagram of patch planting is shown in Fig. 1.

In the next section, we study the scaling of the computational complexity of random instances generated by our PP method by varying both the size of blocks and the number of blocks used for composing the larger instances. For the EA model, we also compare the hardness of planted instances with random systems. More precisely, we use the entropic family size ρ_s [9] to characterize the hardness of the instances.

Classical simulations have been performed by using population annealing while experiments have been performed using the latest D-Wave 2X chip with $N = 1097$ working qubits. The simulation details for the 3-dimensional EA model are shown in Table I. For the Chimera graph, we used all the available qubits and

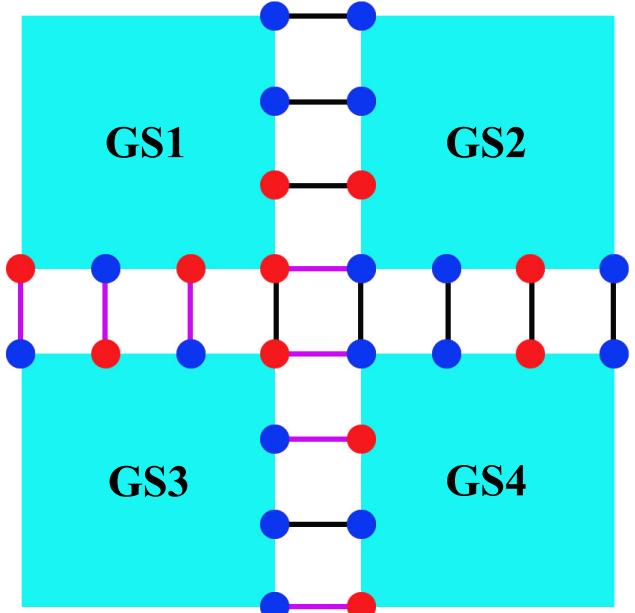


FIG. 1: A schematic diagram of our “patch & planting” method for a 2-dimensional lattice. Firstly, ground states for each block are found assuming free boundary conditions. Then, a ground state of each block is chosen: the surface spins of ± 1 are shown as red and blue circles, respectively. Finally, a coupling is added between neighbor spins of adjacent blocks: a ferromagnetic coupling (black in figure) if the two spins are aligned, or an anti-ferromagnetic if the two-spins are anti-aligned (pink in figure). The direct product of the ground states of the blocks is then the ground state of the larger system.

patched the system using either two, three or four blocks respectively. For the experiments, we used an annealing time of $20 \mu\text{s}$, 100 gauges and 1000 readouts for each gauge. For the numerical simulations, we used a number of replicas (R), number of temperature steps (N_T), number of sweeps (N_S) and temperature (T_0) respectively equal to $R = 2 \times 10^5$, $N_T = 301$, $N_S = 10$ and $T_0 = 0.1$. To compute ρ_0 , we used the following parameters: $R = 10^6$, $N_T = 301$, $N_S = 10$ and $T_0 = 0.1$.

III. RESULTS

A. Correlation between the optimization hardness and the entropic family size ρ_s

The first crucial step in investigating the hardness of instances is to find a good measure which is reliable and yet easy to measure. The success probability of simulated annealing might look a good choice. However, it becomes quickly impossible to compute even for medium size systems. Another possibility consists in looking at

TABLE I: Simulation parameters for the 3-dimensional EA model using population annealing Monte Carlo. Here, L_0 is the block size, L is the linear system size, R is the standard number of replicas, $T_0 = 1/\beta_0$ the lowest temperature simulated, N_T is the number of temperature steps (evenly spaced in β) in the annealing schedule, BC is the type boundary condition [either periodic boundary conditions (PBC) or free boundary conditions (FBC)], and n is the number of disorder realizations studied. For each replica, $N_S = 10$ sweeps have been performed at each temperature. Data for PBC with $L = 8$ are taken from Ref. [8].

L_0	L	R	T_0	N_T	BC	n
4	4	4×10^3	0.2	101	FBC	345600
4	8	10^4	0.2	101	FBC	5000
4	12	5×10^4	0.2	201	FBC	5120
4	16	2×10^5	0.2	301	FBC	1877
4	20	10^6	0.2	401	FBC	194
5	5	10^4	0.2	101	FBC	345600
5	10	10^5	0.2	201	FBC	5000
6	6	2×10^4	0.2	101	FBC	41472
6	12	10^5	0.2	201	FBC	1752
8	8	5×10^4	0.2	201	FBC	23358
8	16	10^6	0.2	301	FBC	624
10	10	10^6	0.2	301	FBC	8000
10	20	2×10^6	0.2	401	FBC	260
8	8	10^5	0.2	101	PBC	5099
12	12	10^6	0.2	201	PBC	3812

very specialized classical algorithms as HFS [34, 35] is for the D-Wave 2X. Nevertheless, the hardness so defined would be based dynamics on the chosen algorithm and may not reflect the structure of the energy landscape. Therefore, in this work we relate the hardness of instances through the entropic family size ρ_s of population annealing. Population annealing (PA) is closely related to simulated annealing (SA) [9] except that it uses a population of replicas and this population is resampled at each temperature step. Unlike SA, PA is designed to simulate the equilibrium Gibbs distribution at each temperature that is traversed. In particular, the resampling step ensures that the population stays close to the equilibrium ensemble.

At low temperature, most of the original population is eliminated in the resampling steps and the final population is descended from a small subset of the initial population. Let n_i be the fraction of population from replica or family i in the initial population, $\rho_s = \lim_{R \rightarrow \infty} R \exp[\sum_i n_i \log n_i]$. Hence, ρ_s represents the characteristic size of survival families: the larger ρ_s is, the more rugged the energy landscape results. Moreover, ρ_s correlates strongly with the integrated autocorrelation time of parallel tempering.

It is well established that ρ_s converges quickly in population size, and it is approximately log-normal distributed

(see Fig. 2 for the analysis of the distribution of Ξ by varying the system size). Therefore, let us define the logarithm of ρ_s

$$\Xi = \log_{10}(\rho_s). \quad (2)$$

Figure 3 shows the correlation between the success probability for SA (with $\beta = 1/T = 5$) and Ξ (data are from Refs. [8, 36]). As expected, the probability of success decreases by increasing Ξ . Indeed, SA struggles more to find the ground state when the energy landscape is more rugged. Therefore, Ξ results the best candidate to measure the hardness of optimization problems.

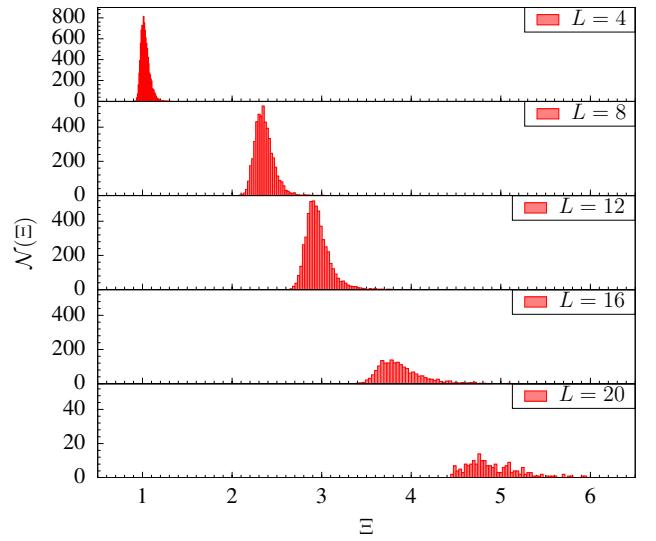


FIG. 2: (Color online) Distribution of Ξ with fixed block size $L_0 = 4$ for various system sizes $L = 4, 8, 12, 16$ and 20 .

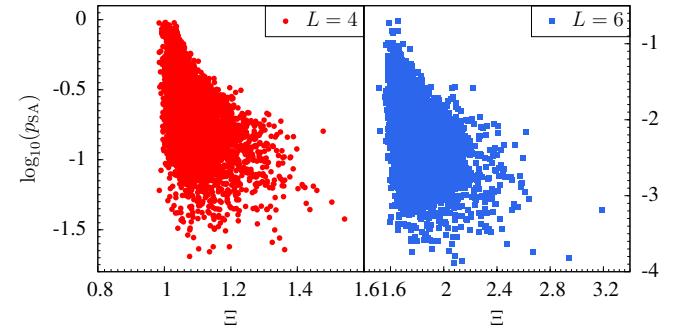


FIG. 3: (Color online) Correlation of the success probability of simulation annealing p_{SA} and the log of entropic family size of population annealing Ξ for $L = 4$ and $L = 6$ at $\beta = 5$ [8, 36]. When Ξ is larger, it is also more difficult to find the ground state. Note that p_{SA} drops very rapidly as L increases, while it is much easier to measure Ξ .

B. Patch & Planting (PP) for the 3D EA model

In this section, we focus on the scaling properties of Ξ for PP instances, by varying either the block sizes L_0 or the system size L . Moreover, we will demonstrate that harder blocks can be used to patch harder instances.

Let us denote with $M = (L/L_0)^3$ the number of blocks of size L_0 . It is well known that, for random instances, ρ_s grows exponentially with L [9]. Since one would expect that having a problem of size L by patching M blocks of size L_0 cannot be harder than a problem composed of a single block of size L , the largest size for the entropic family after patching M blocks is then bounded

$$\Xi(M, L_0) \leq M\Xi(L_0), \quad (3)$$

where $\Xi(M, L_0)$ is the logarithm of the entropy family size after patching M blocks and $\Xi(1, L_0) \equiv \Xi(L_0)$. In Fig. 4, we show the scaling of Ξ by varying the number of blocks M and a power-law fit of the form

$$\Xi(M, L_0) = \Xi(L_0)M^\alpha. \quad (4)$$

As one can see, the price to pay to have instances with an *exactly known* ground state consists in having Ξ scaling sub-linearly with the number of blocks M , with $\alpha = 0.31(3)$. Nevertheless, it also proves that the PP instances become *harder* by increasing the the number of blocks. Therefore, it is guaranteed that, for a sufficient large number of blocks, PP instances are arbitrarily hard. Fig. 5 shows the scaling of the parameter α by increasing the size of blocks L_0 , when the number of blocks is kept fixed to $M = 8$. As one can see from the figure, α remains a constant for a wide range of L_0 implying that α is a characteristic constant for PP problems.

One may also expect to have some benefit by using either larger or harder blocks. Indeed, in both cases, the final effect is to have a larger $\Xi(L_0)$ from which build up the hardness of the PP instances. In Fig. 6 we show the effect of having larger blocks by analyzing the distribution of Ξ at fixed size of the system $L = 16$ using two different block sizes, respectively $L_0 = 4$ (top panel) and $L_0 = 8$ bottom panel. As one can see, PP instances are consistently harder by using larger blocks. Similarly, in Fig. 7 we show the distribution of Ξ by patching PP instances with $M = 8$ blocks of size $L_0 = 6$ by either using *easy* or *hard* blocks. To create the two sets of easy/hard blocks, we have chosen respectively the 8000 easiest and the 8000 harder, among 41472 blocks, and patched together 1000 instances for each of the two sets. Also in this case, PP instances patched with hard blocks are consistently harder than PP instances patched with easy blocks, with the mean values being respectively 3.214(5) and 2.957(4). Therefore, systematic and practical technique developed for engineering hard spin glass in Ref. [37] could be also used in generating hard blocks in patch planting.

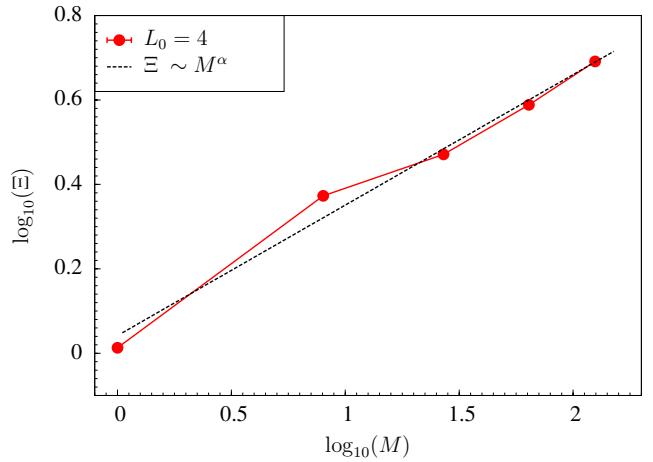


FIG. 4: (Color online) Scaling of the logarithm of the entropic family size $\Xi = \log_{10} \rho_s$ by varying the number of blocks M . The dashed line is a power-law fit of the form $\Xi(L_0)M^\alpha$, where $\Xi(L_0)$ is the entropic family size for a single block. From the fit, we obtain $\alpha = 0.31(3)$.

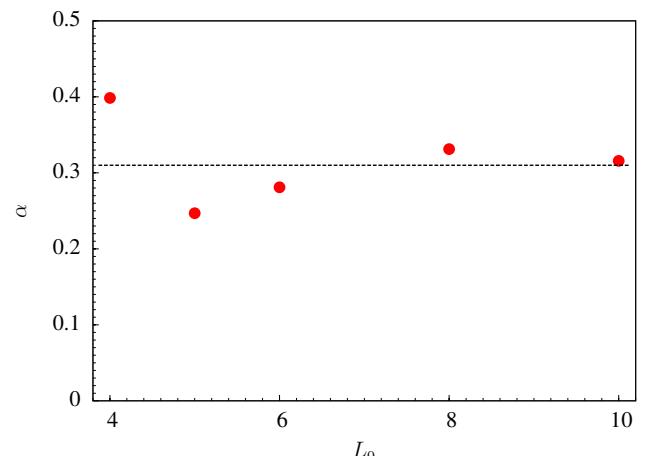


FIG. 5: (Color online) Power-law exponent α for the scaling of $\Xi \sim M^\alpha$, by varying the size of blocks L_0 but keeping fixed the number of blocks to $M = 8$.

Finally, it is interesting to compare the complexity of the patch blocks with random instances. The distribution of Ξ for $L = 8$ and $L = 12$ with different block sizes and random instances are shown in Fig. 8. One can see while the patched instances are generally simpler than random instances, they are not all trivial. There are some overlaps in the distributions. Furthermore, the complexity grows for a fixed system size as the blocks size gets larger.

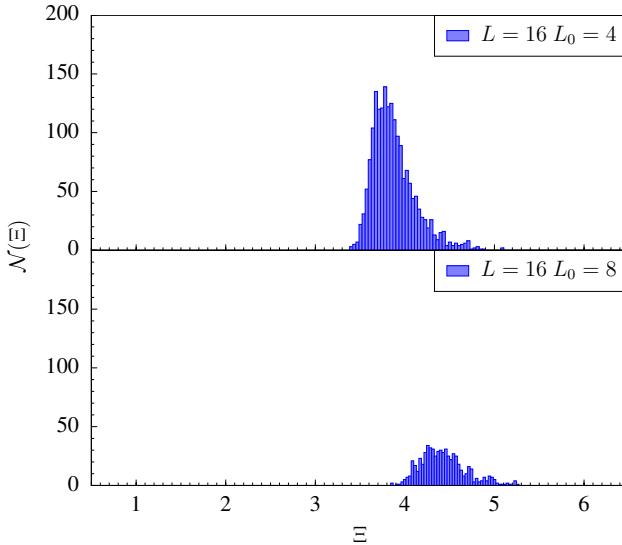


FIG. 6: (Color online) Comparison of the distribution of Ξ for $L = 16$ but with different block sizes $L_0 = 4$ and 8 , respectively. There is a noticeable shift in the distributions of Ξ . Therefore to patch complex instances, one should use as large block sizes as possible.

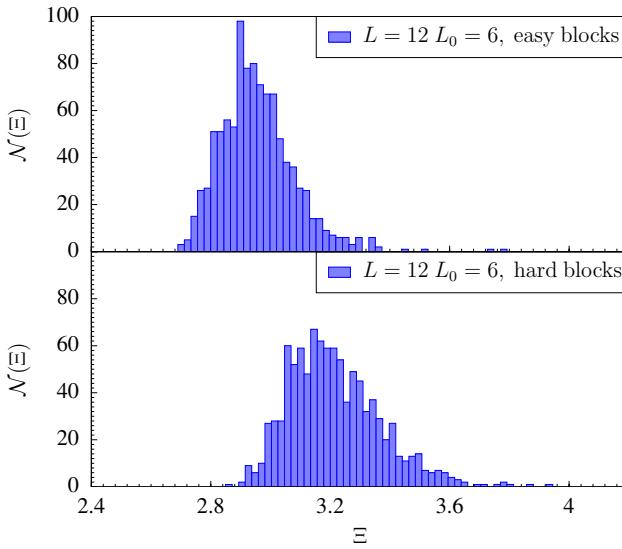


FIG. 7: (Color online) Distributions of Ξ for $M = 8$ blocks of size $L_0 = 6$ using either easy blocks (top panel) or hard blocks (bottom panel). As expected, the distribution shifts to the left when harder blocks are used.

C. The chimera graph

We have also studied patch planting on the chimera graph, for the largest size $N = 1152$ available for the DW2X machine. The purpose is to see the performance of the patch instances for quantum annealing by studied the success probabilities p_{succ} . In addition, the correlation of p_{succ} and Ξ are also checked for this new model.

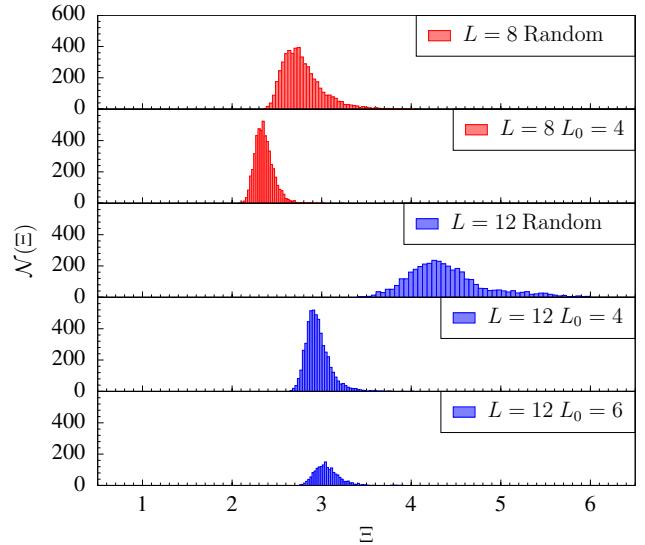


FIG. 8: (Color online) Comparison of the complexity of the patched instances with random instances. The patched instances are generally simpler than random instances as expected, but they are not all trivial. There are some overlaps in the distributions of Ξ . Furthermore, the complexity grows for a fixed system size as the used blocks size gets larger.

Note that there is no spin glass phase on the chimera graph, compared with the 3D case [38]. There are 1097 qubits currently active in DW2X, and we study random systems and patched systems with 2, 3, and 4 blocks that cuts the chimera graph into approximately equal parts horizontally, the direction that cuts few bonds. We studied 1000 instances of each class. For patch planting, the 1000 instances were chosen from 10000 instances, and we have selected the hardest ones.

The distributions of Ξ , p_{succ} and the correlation of the two quantities are shown in Figs. 9, 10 and 11, respectively. With some data mining of only a factor of 10, the instances with two blocks (2B) have approximately the same complexity as the random ones, which are harder than instances with three blocks (3B) and four blocks (4B). Some statistics of the success probabilities are shown in Table II. It is also interesting that the correlation of Ξ and p_{succ} is about the same as the 3D case, even though a spin glass phase is absent on the chimera graph, and the success probability is measured using quantum annealing instead of simulated annealing, showing the complexities are intrinsic to the instances.

IV. SUMMARY

In this paper, we introduced an idea of patch planting for creating arbitrary large and nontrivial spin glass instances with known ground states. We studied in detail of the scaling of the complexity of the patched instances, and compared them with random instances using popu-

TABLE II: Some statistics of the DW2X success probability p_{succ} in Fig. 10 for random instances and patched instances with two blocks (2B), three blocks (3B) and four blocks (4B). There are 1000 instances each, the the patched instances were chosen from the hardest ones out of 10000 instances in each class. p_{\min} , p_{\max} and p_{ave} are the minimum, maximum and average values of p_{succ} and f is the fraction of instances with $p_{\text{succ}} = 0$.

Statistic	Random	2B	3B	4B
p_{\min}	0	0	0	0
p_{\max}	2.55×10^{-3}	1.84×10^{-3}	3.49×10^{-2}	1.76×10^{-1}
p_{ave}	2.04×10^{-5}	1.29×10^{-5}	1.15×10^{-3}	1.27×10^{-2}
f	0.831	0.775	0.185	0.011

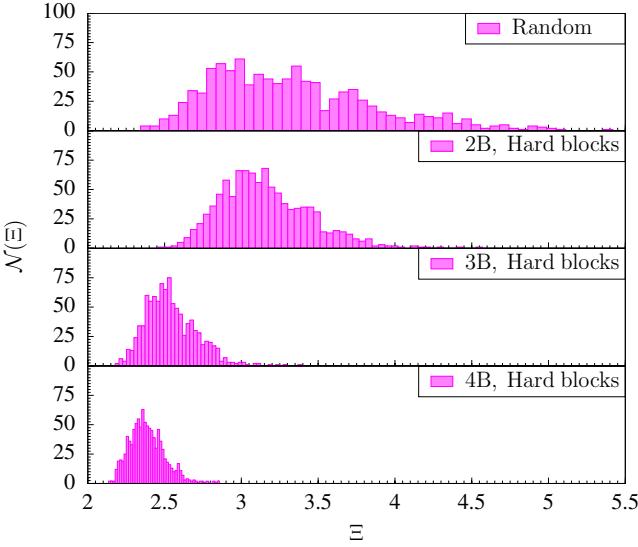


FIG. 9: (Color online) The distributions of Ξ on the chimera graph for $N = 1152$ for random instances and patched instances with two blocks (2B), three blocks (3B) and four blocks (4B). There are 1000 instances each, and the patched instances were chosen from the hardest ones out of 10000 instances in each class. Note that 2B and Random have about the same complexity.

lution annealing and the DW2X machine. Furthermore, we have shown that to patch hard instances, one should use as large blocks as possible and also as hard blocks as possible. Our method is easy to use and could be used as benchmark instances for future generation quantum annealing machines, and classical as well as quantum mechanical optimization algorithms.

Acknowledgments

W.W. and H.G.K. acknowledge support from the National Science Foundation (Grant No. DMR-1151387). We thank Jon Machta and Ethan Brown for helpful discussions. The research of H.G.K. is based upon work

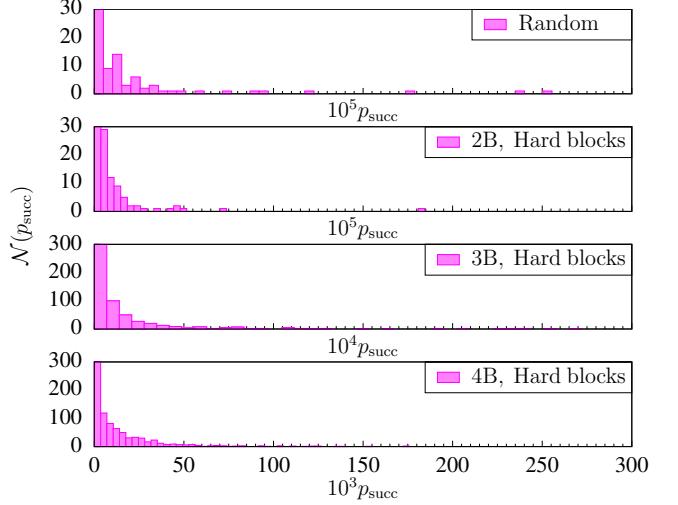


FIG. 10: (Color online) The distributions of the success probability p_{succ} on the chimera graph for $N = 1152$ for random instances, 2B, 3B and 4B. There are 1000 instances each, and the patched instances were chosen from the hardest ones out of 10000 instances in each class. Note that 2B and Random have about the same complexity.

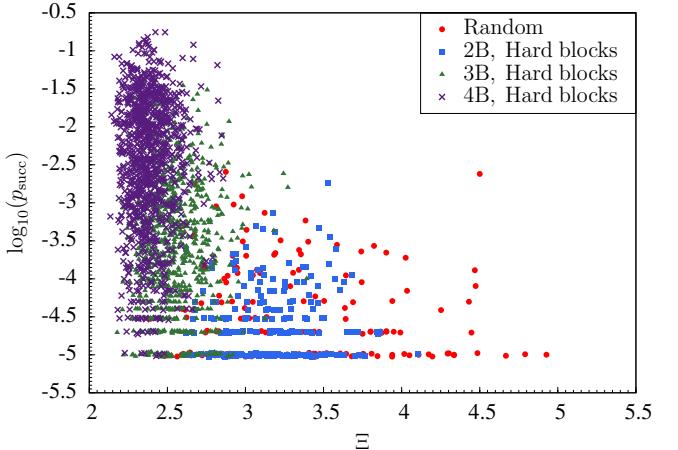


FIG. 11: (Color online) Correlation of Ξ and p_{succ} on the chimera graph for $N = 1152$ for random instances, 2B, 3B and 4B. Note that the feather is essentially the same as Fig. 3, even though there is no spin glass phase on the chimera graph and the success probability is measured using quantum annealing instead of simulated annealing, showing the complexities are intrinsic to the instances.

supported in part by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via MIT Lincoln Laboratory Air Force Contract No. FA8721-05-C-0002. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of ODNI, IARPA, or the U.S. Gov-

ernment. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purpose notwithstanding any copyright annotation thereon. We

thank Texas A&M University for access to their Ada and Curie clusters.

[1] S. F. Edwards and P. W. Anderson, *Theory of spin glasses*, J. Phys. F: Met. Phys. **5**, 965 (1975).

[2] S. Kirkpatrick, C. D. Gelatt, Jr., and M. P. Vecchi, *Optimization by simulated annealing*, Science **220**, 671 (1983).

[3] R. H. Swendsen and J.-S. Wang, *Replica Monte Carlo simulation of spin-glasses*, Phys. Rev. Lett. **57**, 2607 (1986).

[4] C. Geyer, in *23rd Symposium on the Interface*, edited by E. M. Keramidas (Interface Foundation, Fairfax Station, VA, 1991), p. 156.

[5] K. Hukushima and K. Nemoto, *Exchange Monte Carlo method and application to spin glass simulations*, J. Phys. Soc. Jpn. **65**, 1604 (1996).

[6] K. Hukushima and Y. Iba, in *The Monte Carlo method in the physical sciences: celebrating the 50th anniversary of the Metropolis algorithm*, edited by J. E. Gubernatis (AIP, 2003), vol. 690, p. 200.

[7] E. Zhou and X. Chen, in *Proceedings of the 2010 Winter Simulation Conference (WSC)* (2010), p. 1211.

[8] W. Wang, J. Machta, and H. G. Katzgraber, *Comparing Monte Carlo methods for finding ground states of Ising spin glasses: Population annealing, simulated annealing, and parallel tempering*, Phys. Rev. E **92**, 013303 (2015).

[9] W. Wang, J. Machta, and H. G. Katzgraber, *Population annealing: Theory and application in spin glasses*, Phys. Rev. E **92**, 063307 (2015).

[10] A. K. Hartmann and H. Rieger, *Optimization Algorithms in Physics* (Wiley-VCH, Berlin, 2001).

[11] A. K. Hartmann and H. Rieger, *New Optimization Algorithms in Physics* (Wiley-VCH, Berlin, 2004).

[12] C. De Simone, M. Diehl, M. Jünger, P. Mutzel, G. Reinelt, and G. Rinaldi, *Exact ground states in spin glasses: New experimental results with a branch-and-cut algorithm*, J. Stat. Phys. **80**, 487 (1995).

[13] J. Houdayer, *A cluster Monte Carlo algorithm for 2-dimensional spin glasses*, Eur. Phys. J. B. **22**, 479 (2001).

[14] Z. Zhu, A. J. Ochoa, and H. G. Katzgraber, *Efficient Cluster Algorithm for Spin Glasses in Any Space Dimension*, Phys. Rev. Lett. **115**, 077201 (2015).

[15] T. Kadowaki and H. Nishimori, *Quantum annealing in the transverse Ising model*, Phys. Rev. E **58**, 5355 (1998).

[16] H. Nishimori, *Statistical Physics of Spin Glasses and Information Processing: An Introduction* (Oxford University Press, New York, 2001).

[17] E. Farhi, J. Goldstone, S. Gutmann, J. Lapan, A. Lundgren, and D. Preda, *A quantum adiabatic evolution algorithm applied to random instances of an NP-complete problem*, Science **292**, 472 (2001).

[18] G. Santoro, E. Martoňák, R. Tosatti, and R. Car, *Theory of quantum annealing of an Ising spin glass*, Science **295**, 2427 (2002).

[19] J. Roland and N. J. Cerf, *Quantum search by local adiabatic evolution*, Phys. Rev. A **65**, 042308 (2002).

[20] S. Boixo, T. Albash, F. M. Spedalieri, N. Chancellor, and D. A. Lidar, *Experimental signature of programmable quantum annealing*, Nat. Commun. **4**, 2067 (2013).

[21] S. Boixo, T. F. Rønnow, S. V. Isakov, Z. Wang, D. Wecker, D. A. Lidar, J. M. Martinis, and M. Troyer, *Evidence for quantum annealing with more than one hundred qubits*, Nat. Phys. **10**, 218 (2014).

[22] T. F. Rønnow, Z. Wang, J. Job, S. Boixo, S. V. Isakov, D. Wecker, J. M. Martinis, D. A. Lidar, and M. Troyer, *Defining and detecting quantum speedup*, Science **345**, 420 (2014).

[23] S. Boixo, V. N. Smelyanskiy, A. Shabani, S. V. Isakov, M. Dykman, V. S. Denchev, M. H. Amin, A. Y. Smirnov, M. Mohseni, and H. Neven, *Computational multiqubit tunnelling in programmable quantum annealers*, Nat. Comm. **7**, 10327 (2016).

[24] D. A. Lidar and T. A. Brun, eds., *Quantum Error Correction* (Cambridge University Press, Cambridge, UK, 2013).

[25] K. L. Pudenz, T. Albash, and D. A. Lidar, *Quantum annealing correction for random Ising problems* (2014), arXiv:quant-ph/1408.4382.

[26] A. Perdomo-Ortiz, B. O’Gorman, J. Fluegmann, R. Biswas, and V. N. Smelyanskiy, *Determination and correction of persistent biases in quantum annealers* (2015), (arXiv:quant-ph/1503.05679).

[27] W. Vinci, T. Albash, G. Paz-Silva, I. Hen, and D. A. Lidar, *Quantum annealing correction with minor embedding*, Phys. Rev. A **92**, 042310 (2015).

[28] S. Mandrà, G. G. Guerreschi, and A. Aspuru-Guzik, *Adiabatic quantum optimization in the presence of discrete noise: Reducing the problem dimensionality*, Phys. Rev. A **92**, 062320 (2015).

[29] D. Venturelli, S. Mandrà, S. Knysh, B. O’Gorman, R. Biswas, and V. Smelyanskiy, *Quantum Optimization of Fully Connected Spin Glasses*, Phys. Rev. X **5**, 031040 (2015).

[30] M. W. Johnson, M. H. S. Amin, S. Gildert, T. Lanting, F. Hamze, N. Dickson, R. Harris, A. J. Berkley, J. Johansson, P. Bunyk, et al., *Quantum annealing with manufactured spins*, Nature **473**, 194 (2011).

[31] B. Heim, T. F. Rønnow, S. V. Isakov, and M. Troyer, *Quantum versus classical annealing of ising spin glasses*, Science **348**, 215 (2015).

[32] S. Mandrà, Z. Zhu, W. Wang, A. Perdomo-Ortiz, and H. G. Katzgraber, *Strengths and Weaknesses of Weak-Strong Cluster Problems: A Detailed Overview of State-of-the-art Classical Heuristics vs Quantum Approaches*, Physical Review A **94**, 022337 (2016), (arXiv:1604.01746).

[33] S. Mandrà, Z. Zhu, and H. G. Katzgraber, *Exponentially-biased ground-state sampling of quantum annealing machines with transverse-field driving hamiltonians*, arXiv preprint arXiv:1606.07146 (2016).

[34] F. Hamze and N. de Freitas, in *Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence* (AUAI Press, Arlington, Virginia, United States, 2004), UAI ’04, p. 243, ISBN 0-9749039-0-6.

[35] A. Selby, *Efficient subgraph-based sampling of Ising-*

type models with frustration (2014), (arXiv:cond-mat/1409.3934).

[36] W. Wang, J. Machta, and H. G. Katzgraber, *Evidence against a mean-field description of short-range spin glasses revealed through thermal boundary conditions*, Phys. Rev. B **90**, 184412 (2014).

[37] J. Marshall, V. Martin-Mayor, and I. Hen, *Practical engineering of hard spin-glass instances*, Phys. Rev. A **94**, 012320 (2016).

[38] H. G. Katzgraber, F. Hamze, and R. S. Andrist, *Glassy Chimeras Could Be Blind to Quantum Speedup: Designing Better Benchmarks for Quantum Annealing Machines*, Phys. Rev. X **4**, 021008 (2014).